

Introduction of IHEP Monitoring System

Qingbao Hu

hugb@ihep.ac.cn

Computing Center, IHEP

Chinese Academy of Sciences

ISGC 2016 Monitoring BoF

Status of IHEP Cluster

- ~ 1,122 work nodes
 - ~ 13,500 CPU cores
- ~ 5PB disk storage
 - Lustre, gLuster, openAFS, etc.
- ~ 5PB tape storage
 - Two IBM 3584 tape libraries, LTO4 tape
 - Modified CERN CASTOR 1.7



Cluster built
with blades

Tape libraries



Monitoring System Overview

- System overview

Monitoring system of IHEP

Ganglia


Recording the performance of different resource groups

Icinga

Monitoring the status of cluster devices and services

Logger Analysis

Collecting more comprehensive data & providing an overview of the whole cluster health status



Ganglia

- **Monitoring the health of the cluster**
 - System load
 - CPU utilization
 - Network bandwidth and traffic
 - Memory usage
- **Usage**
 - Records history status of the cluster
 - Helps system manager to fix problem



Icinga

- **Created as a fork of the Nagios**
 - Plug-in design
 - Active check of the service
- **Polling agents we developed**
 - More services monitored
 - Some crashed service faults can be recovered automatically
 - Alarms sent to system manager on time (email, SMS)



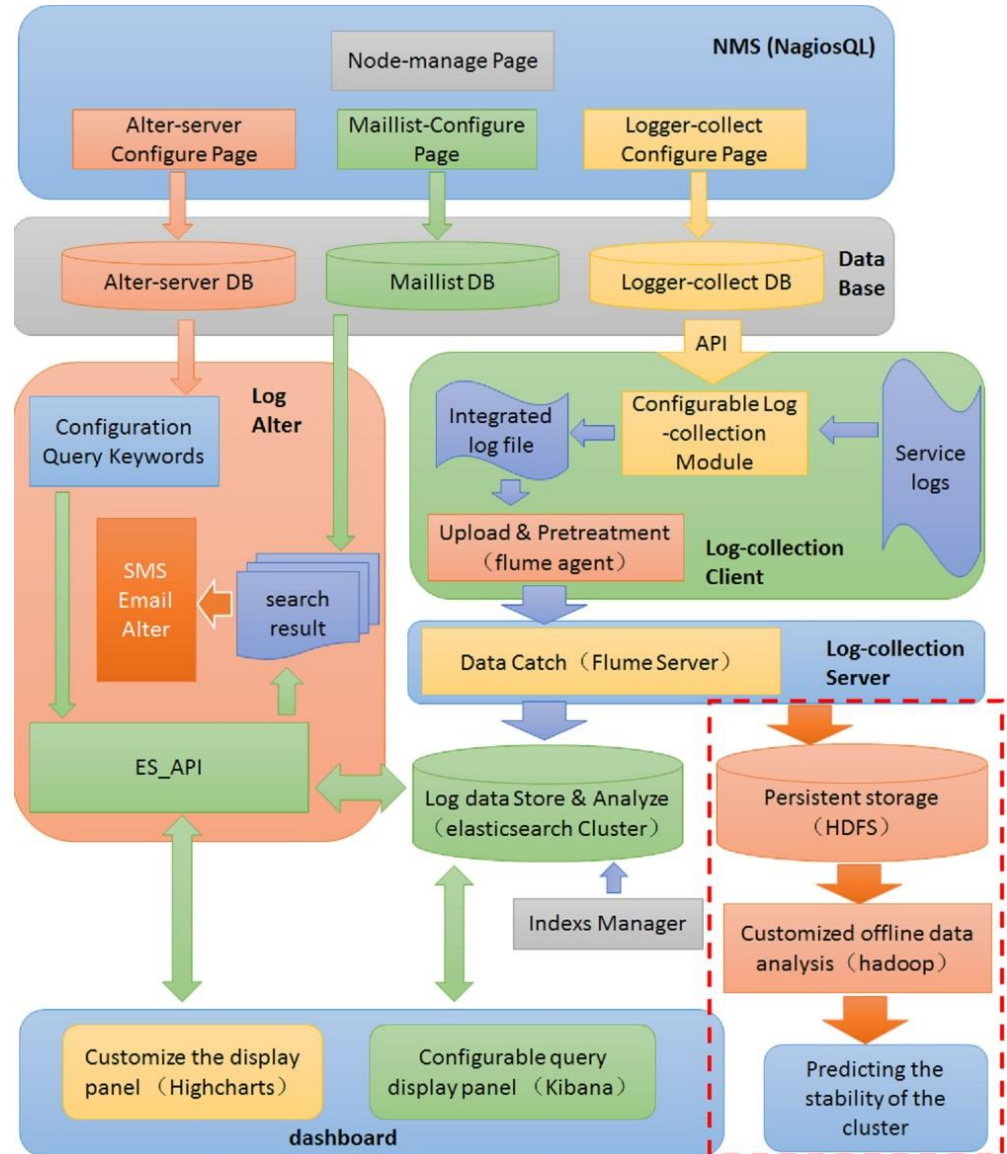
Logger-analysis Monitoring

- **Provides a novel monitoring based on log analysis**
- **Flume**
 - A distributed, reliable, and available service for efficiently collecting, aggregating, and moving large amounts of log data. It has a simple and flexible architecture based on streaming data flows.
 - It is robust and fault tolerant with tunable reliability mechanisms and many failover and recovery mechanisms.
- **Elasticsearch**
 - Distributed, scalable, and highly available
 - Real-time search and analytics capabilities
 - RESTful API




Logger-analysis Monitoring

- IHEPlogger Architecture




Logger-analysis Monitoring

- **Flume + Elasticsearch**
 - Performance optimized
 - All logs are pre-processed before collected by flume
 - » No log lost
 - » Easy to be queried in dedicated groups



Next plan

- **Logger analysis**
 - Archiving of the log data
 - Store archive data on HDFS
 - Offline log mining based on Hadoop or Storm
- **Integrate remote sites monitoring**



Thank you!
Any Question?

